

PATENT

Atty. Docket No.: ROC920030215US1 (IBM/264)

Confirmation No. 2241

CERTIFICATE OF ELECTRONIC TRANSMISSION

I hereby certify that this correspondence for Application No. 10/691,296 is being electronically transmitted to Technology Center 2166, via EFS-WEB, on December 6, 2006.

/Scott A. Stinebruner/
Scott A. Stinebruner, Reg. No. 38,323

December 6, 2006
Date

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant: Paul Reuben Day et al. Art Unit: 2166
Application No.: 10/691,296 Examiner: Emeka Ebirim
Filed: October 22, 2003
For: OPTIMIZATION OF QUERIES USING RETRIEVAL STATUS OF
RESOURCES USED THEREBY

RESPONSE AFTER NON-FINAL REJECTION

Mail Stop AMENDMENT
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

This paper is submitted in reply to the Office Action dated September 6, 2006, within the three-month period for response. Reconsideration and allowance of all pending claims are respectfully requested.

In the subject Office Action, claim 29 was rejected under 35 U.S.C. § 101 and under 35 U.S.C. § 112 first paragraph. Moreover, claims 3-11, 16-24 and 29 were rejected under 35 U.S.C. § 102(e) as being anticipated by U.S. Patent Application Publication No. 2004/0073549 by Turkel et al. In addition, claims 1-2, 12-15 and 25-28 were rejected under 35 U.S.C. § 103(a) as being unpatentable over U.S. Patent No. 6,266,658 to Adya et al. in view of Turkel et al.

Applicants respectfully traverse the Examiner's rejections to the extent that they are maintained.

Now turning to the subject Office Action and initially to the § 101 rejection of claim 29, the Examiner will note that Applicants have previously amended the claim to

recite a "tangible computer readable signal bearing medium." The Examiner asserts that claim 29 is not limited to tangible embodiments in view of the disclosure, but considering the fact that the claim itself is explicitly limited to a tangible medium, Applicants fail to comprehend the Examiner's basis for the continued rejection of claim 29. The claim recites a tangible medium, and as such, Applicants submit that the claim is fully compliant with the Office's current interpretation of § 101. In addition, it is Applicants' understanding that it is acceptable practice under the Guidelines for the Examiner to suggest to Applicants what amendments could be made to a claim rejected under §101 in order to place the claim in statutory condition. As such, Applicants respectfully request that if the Examiner chooses to continue the rejection of claim 29, the Examiner provide Applicants with a proposed amendment to the claim in the Examiner's next action. Reconsideration and withdrawal of the § 101 rejection are therefore respectfully requested.

Claim 29 is also apparently rejected under 35 U.S.C. § 112, first paragraph for failing to be supported by a "credible asserted utility." The claim, however, explicitly recites "optimiz[ing] a database query," which is undoubtedly a credible utility. The Examiner asserts that one of ordinary skill in the art would not know how to use the claimed invention. However, in view of the description of a program product and a signal bearing medium at page 10 of the Application as filed, as well as the functionality of the claimed program code as described at pages 12-14 and shown in Figs. 3-4, Applicants submit that one of ordinary skill in the art would readily appreciate how to use the claimed invention. Reconsideration and withdrawal of the § 112 rejection are therefore respectfully requested.

Now turning to the art-based rejections, and initially to the § 102 rejection of independent claim 3, this claim generally recites a method of optimizing a database query, which includes determining a retrieval status for a resource used by the database query, and generating an access plan for the database query using the determined retrieval status for the resource.

As Applicants have mentioned previously, Applicants' invention addresses a problem experienced by conventional query optimizers where the estimated cost of an

access plan as calculated by a query optimizer may not accurately represent the actual cost for executing an access plan based upon a given runtime environment, in particular due to inaccurate estimations of input/output costs that are based upon the amount of time required to retrieve resources necessary for executing a query. Whenever a cost estimate is incorrect, a risk exists that a query optimizer will not select the optimal access plan, and thus a query will not execute with optimal efficiency.

As an example, a resource such as an index often is able to substantially accelerate the execution of a query. However, if an index is stored in persistent storage, there is an input/output cost associated with retrieving the index into working memory (e.g., RAM) before that index can be used during execution of the query. In some instances, the cost to load the index into working memory may be too excessive, and may result in the selection of another access plan that does not use the index. On the other hand, if the index has already been loaded into working memory before the query is executed (e.g., if a prior query has already loaded the index), the input/output cost associated with loading the working memory is practically eliminated, and may make an access plan that uses the index more efficient than competing access plans. Conventional query optimizers, however, are unable to ascertain whether an index has or has not already been loaded into a working memory when selecting an access plan for a query that might use such an index, and as such, are often required to assume that an index will be required to be loaded. As such, if a particular access plan that uses the index would be more efficient than competing access plans if the cost of loading the index was not taken into account, that access plan will typically not be selected even if the index was already loaded into working memory when the query is executed.

Embodiments of Applicants' invention, on the other hand, are capable of using a "retrieval status" of a resource such as a database table, file, index, or hash table to generate a more accurate input/output cost estimate for an access plan, and thus enable a query optimizer to select an optimal access plan given the current runtime conditions under which the access plan will execute.

Claim 3 is rejected based upon Turkel, citing Figs. 2 and 5 and paragraphs [0087], [0092]-[0093] and [0127] thereof. These passages disclose the use of caches 34 in

execution blocks 26 of a storage access system 20. The execution blocks provide query interfaces that are accessed by clients, and a sweeping unit 22 is used to supply the execution blocks 26 with data from the database to service client queries (paragraphs [0081-0085]). As discussed at paragraphs [0090]-[0093], data may be swapped in and out of caches 34 based upon frequency of use, so data that is frequently used by an execution block may be retained in its cache to speed access to the data.

Fig. 2 of Turkel and the accompanying disclosure at paragraphs [0092]-[0093] arguably disclose determining the status of a requested resource, e.g., whether requested data is currently in a cache. However, claim 3 specifically recites "generating an access plan for [a] database query using the determined retrieval status for [a] resource," a feature that is not disclosed or suggested by Turkel.

Specifically, the Examiner relies on Fig. 5 and paragraphs [0092] and [0127] of Turkel for allegedly disclosing this feature. However, these passages do not disclose any generation of an access plan that uses the determined retrieval status of a resource. The passage at paragraph [0092], for example, discusses in connection with Fig. 2 a storage management system that handles data requests; however, these data requests are distinguished from queries, and appear to be requests for data that may be issued in connection with executing a query elsewhere in a database management system. Nonetheless, even if these data requests were queries, the passage does not disclose or suggest generating such requests based upon a retrieval status of a resource.

The passage at paragraph [0127], which discusses Figs. 4 and 5, discloses that a preprocessor 122 may provide a data transfer unit 132 with a list of attributes required for a particular group of queries, presumably so that the data transfer unit can retrieve the appropriate data necessary to resolve that group of queries. Of note, however, at no point is the execution plan of a query (which is generated in block 156 of Fig. 5, as described at paragraphs [0115]-[0116]) ever generated based upon the retrieval status of any of the data required to resolve a query. The generation of the execution plan in Turkel (performed by a preprocessing unit 122) is not conditioned on the retrieval status of any data or resource in the database. At the most, Turkel discloses that a preprocessing unit 122 can notify a data transfer unit 132 to retrieve data; however, there is no disclosure

anywhere in the reference of the data transfer unit 132 (or any other retrieval component, e.g., cache 140 or controller 192) notifying a preprocessing unit 122 of the status of any data required for a query. In short, the retrieval status of any resource is never used by a preprocessing unit to decide how to generate an access or execution plan.

The Examiner may also note that Turkel discloses at paragraphs [0130]-[0131] that a post processing unit 136 may query a cache 140. However, the post processing unit does not generate any access or execution plan, so the fact that the cache may be queried still falls short of disclosing basing the generation of an access plan on the retrieval status of a resource.

As such, Turkel does not disclose the recited feature of generating an access plan based upon a retrieval status of a resource. Claim 3 is therefore novel over Turkel.

Claim 3 is also non-obvious over Turkel and the other prior art of record, most notably Adya. In particular, Turkel and the other art of record fail to suggest the desirability of using the retrieval status of a resource in connection with generating an access plan. In fact, Turkel appears to teach away from such an approach given that the reference attempts to control a data transfer unit to retrieve data as is needed by a particular query, which in many respects is an opposite approach to that undertaken in the invention recited in claim 3. Put into the language of claim 3, Turkel, if anything, attempts to adapt the retrieval status of a resource to a given access plan, while claim 3 effectively attempts to adapt an access plan to a given retrieval status of a resource.

Applicants therefore respectfully submit that claim 3 is novel and non-obvious over Turkel and the other prior art of record. Reconsideration and allowance of independent claim 3, and of claims 4-15 which depend therefrom, are therefore respectfully requested.

Next turning to independent claims 16 and 29, each of these claims recite in part program code configured to a database query by "determining a retrieval status for a resource used by the database query, and generating an access plan for the database query using the determined retrieval status for the resource." As discussed above in connection with claim 3, this combination of features is not disclosed or suggested by Turkel. Accordingly, claims 16 and 29 are novel and non-obvious over Turkel for the same

reasons as presented above for claim 3. Reconsideration and allowance of independent claims 16 and 29, as well as of claims 17-28 which depend therefrom, are therefore respectfully requested.

Next, with respect to the § 103 rejection of claim 1, this claim generally recites a method of performing a database query. The method includes generating an access plan for the database query, estimating a percentage of the resource that is currently resident in working memory, estimating a cost for the access plan using the estimated percentage, and selectively executing the access plan based upon the estimated cost. The access plan uses at least one resource capable of being retrieved into working memory, where the resource is selected from the group consisting of a database file, a database table, an index, a temporary result set, a temporary file, and a hash table.

In rejecting claim 1, the Examiner primarily relies on Adya, and in particular col. 1, lines 22-24, col. 3, line 40, col. 7, lines 1-7, col. 8, lines 25-30 and col. 10, lines 11-14. However, in none of these passages, nor anywhere else in Adya, can Applicants find any disclosure of a number of features recited in claim 1, e.g., "estimating a percentage of [a] resource that is current resident in working memory" and "estimating a cost for [an] access plan using the estimated percentage."

Instead, Adya is directed to a methodology for determining a set of indexes to be proposed to an administrator for the purpose of optimizing a runtime environment for a given workload. The methodology attempts to determine what indexes should be generated or maintained in storage based upon limited available storage. The goal of Adya is to recommend a set of indexes that will have the greatest positive impact on performance within the limited amount of storage space that can hold those indexes.

The specific passages relied upon by the Examiner disclose, at the most, using cost estimates for queries to rank potential indexes for recommendation to a systems administrator for a DBMS. The passages do not, in fact, even deal with executing queries; instead they deal with recommending indexes that a system administrator could manually generate or update to optimize system workload in the future. Indeed, the cost calculation disclosed at col. 8, lines 25-30 is not a cost of an access plan, it is the

estimated cost of an index to determine whether that index, if created, would substantially improve the performance of the system.

Adya furthermore does not ever attempt to determine whether or not an index even exists in any form of storage, much less doing so for the purpose of calculating the cost of an access plan that might use such an index. More importantly, Adya does not disclose determining a retrieval status of an index, and using that retrieval status to adjust the cost estimate for an access plan.

Given that Adya does not disclose determining a retrieval status of an index, Applicants submit that Adya cannot be interpreted as disclosing the estimation of a percentage of an index (or any other required resource) that is resident in a working memory, or the estimation of a cost for an access plan using any estimated percentage.

The Examiner appears to acknowledge these shortcomings in Adya, and relies on Turkel, and in particular Figs. 2 and 5 and paragraphs [0006]-[0007] and [0092]-[0094] thereof. However, as discussed above in connection with claim 3, Turkel fails to disclose or suggest generating an access plan based upon the retrieval status of a resource. Likewise, Turkel fails to disclose or suggest estimating a cost for an access plan based upon the retrieval status of a resource. The cited passages at paragraphs [0092]-[0094] disclose, at the most, how a storage access system can determine what data to place in a cache and what data to place on a bus. At no point is this information ever used either in the generation of an access plan or the generation of a cost for an access plan. Indeed, the only discussion of an "execution plan" in Turkel does not address costing competitive plans or deciding whether or not to execute certain plans based upon cost. The cited passages at paragraphs [0006]-[0007] are merely relied upon for motivation, and similarly do not disclose any such concepts.

Furthermore, claim 1 specifically recites estimating the percentage of a resource that is resident in a working memory. Turkel is completely silent with respect to estimating the percentage of anything, much less that of a resource used by an access plan. The string "percent" is not even present anywhere in the reference. Further, as noted above, Adya similarly does not disclose estimating the percentage of a resource that is resident in a working memory.

Accordingly, Applicants submit that the proposed combination of Adya and Turkel does not disclose or suggest each and every feature of claim 1, and as such, the Examiner has failed to establish a *prima facie* case of obviousness as to claim 1. Furthermore, given that neither reference suggests the estimation of the cost of an access plan based upon the retrieval status of a resource, including but not limited to the percentage of that resource that may be resident in a working memory, Applicants submit that the Examiner has failed to provide objective evidence of a motivation to modify the proposed combination to incorporate such features. Reconsideration and allowance of claim 1, and of claim 2 which depends therefrom, are therefore respectfully requested.

As a final matter, Applicants traverse the Examiner's rejections of the dependent claims based upon the dependency of these claims on the aforementioned independent claims. Applicants do wish to point out, however, that a number of these claims additionally recite features that further distinguish the claims from the art of record. For example, with respect to claims 8, 13 and 21, which recite in part the determination of a percentage of a resource that has been retrieved, the cited passages are completely silent with respect to determining a percentage of a resource that has been retrieved. The Examiner has therefore failed to establish that this feature is disclosed by Turkel or Adya.

With respect to claims 10-11 and 23-24, which recite in part the maintenance of a percentage of a resource that is working memory with a resource manager, the Examiner relies on Turkel, and in particular Paragraphs [0093] and [0087]. However, as noted above, there is no disclosure in the reference related to determining the percentage of a resource. The Examiner's rejection therefore falls far short of disclosing or suggesting each and every feature of these claims.

With respect to claims 14 and 27, which recite in part the selective generation of an access plan based upon the comparison of a current retrieval status with a stored assumption for a stored access plan, the cited passages are also completely silent with respect to this concept. The Examiner has therefore failed to establish that this feature is disclosed by Turkel or Adya.

Finally, with respect to claims 15 and 28, which recite the determination of whether a beginning portion of a resource is resident in working memory, coupled with

weighting a cost based upon this determination, the cited passages are completely silent with respect to different portions of a resource, much less performing any weighting on this basis. The Examiner has therefore failed to establish that these features are disclosed by Turkel or Adya.

In summary, Applicants respectfully submit that all pending claims are novel and non-obvious over the prior art of record. Reconsideration and allowance of all pending claims are therefore respectfully requested. If the Examiner has any questions regarding the foregoing, or which might otherwise further this case onto allowance, the Examiner may contact the undersigned at (513) 241-2324. Moreover, if any other charges or credits are necessary to complete this communication, please apply them to Deposit Account 23-3000.

Respectfully submitted,

December 6, 2006

Date

/Scott A. Stinebruner/

Scott A. Stinebruner
Reg. No. 38,323
WOOD, HERRON & EVANS, L.L.P.
2700 Carew Tower
441 Vine Street
Cincinnati, Ohio 45202
Telephone: (513) 241-2324
Facsimile: (513) 241-6234